

Predictive Performance Comparison of Robust Classifiers on ε -Contaminated High Dimension Low Sample Size Data

Necla Gündüz^{*,†}, Ernest Fokoué[‡]

[†]*Department of Statistics, Faculty of Science
Gazi University, Ankara, Turkey
e-mail: ngunduz@gazi.edu.tr*

[‡]*School of Mathematical Sciences
Rochester Institute of Technology
98 Lomb Memorial Drive, Rochester, NY 14623, USA
e-mail: epfeqa@rit.edu*

Abstract: The robustification of pattern recognition techniques has been the subject of intense research in recent years. Despite the multiplicity of papers on the subject, very few articles have deeply explored the topic of robust classification in the high dimension low sample size context. In this work, we explore and compare the predictive performances of robust classification techniques with a special concentration on robust discriminant analysis and robust PCA applied to a wide variety of large p small n data sets. We also explore the performance of random forest by way of comparing and contrasting the differences single model methods and ensemble methods in this context. Our work reveals that Random Forest, although not inherently designed to be robust to outliers, substantially outperforms the existing techniques specifically designed to achieve robustness. Indeed, random forest emerges as the best predictively on both real life and simulated data.

AMS 2000 subject classifications: Primary 60K35; secondary 60K35.

Keywords and phrases: High-dimensional, Robust, Prediction Error, Contamination, Outlier, Large p small n , Projection Pursuit, Discriminant Analysis, Random Forest, Minimum Covariance Determinant.

1. Introduction

We are given a data set $\mathcal{D} = \{(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_n, \mathbf{y}_n)\}$ where $\mathbf{x}_i \in \mathbb{R}^{p \times 1}$ with $n \lll p$ and $\mathbf{y}_i \in \{1, 2, \dots, K\}$, and we consider the task of building a classifier $\hat{f}(\cdot)$ such that the true error is as small as possible. In this context, we shall use the average test error $\text{AVTE}(\cdot)$, as our measure of predictive performance, namely

$$\text{AVTE}(\hat{f}) = \frac{1}{R} \sum_{r=1}^R \left\{ \frac{1}{m} \sum_{j=1}^m \ell(\mathbf{y}_j^{(r)}, \hat{f}^{(r)}(\mathbf{x}_j^{(r)})) \right\}, \quad (1.1)$$

where $(\mathbf{x}_j^{(r)}, \mathbf{y}_j^{(r)})$ is the j th observation from the test set at the r th random replication of the split of the data. Here, we use the zero-one loss defined by

$$\ell(\mathbf{y}_j^{(r)}, \hat{f}^{(r)}(\mathbf{x}_j^{(r)})) = 1_{\{\mathbf{y}_j^{(r)} \neq \hat{f}^{(r)}(\mathbf{x}_j^{(r)})\}} = \begin{cases} 1 & \text{if } \mathbf{y}_j^{(r)} \neq \hat{f}^{(r)}(\mathbf{x}_j^{(r)}) \\ 0 & \text{otherwise.} \end{cases} \quad (1.2)$$

The pattern recognition literature is filled with techniques created and developed to solve precisely this problem. Amongst others, logistic regression, discriminant analysis, k -nearest neighbors, classification trees, support vector machine, just a name of few. In this paper, we consider data sets of very special kind, namely the so-called High Dimension Low Sample Size (HDLSS). Also known as large p , small n since for this type of data, $n \lll p$, i.e., n is much less than p . Data sets of this type are very common these days especially from the fields of study involving microarray gene expression used in diagnosing and helping cure disease such as cancer. As a matter fact, we consider five such data sets in this paper containing information about various forms of cancer, namely prostate, lymphoma, lung, colon, leukemia, brain. Traditional classification techniques like logistic regression,

*Corresponding author

discriminant analysis and k -nearest neighbors fail miserably on this kind of data, mainly due to the fact that the condition $n \lll p$ leads to illposedness and thereby the inability of those methods to even have a solution. In the case of k -nearest neighbors for instance, the $n \lll p$ condition leads to a severe case of curse of dimensionality, since the concept of neighbor then becomes loose and ill-defined when the dimension of the space is far larger than the number of observations available [Kondo et al. \(2012\)](#). Several approaches have been proposed to achieve optimal classification in this HDLSS context. One of the earliest is regularized discriminant analysis proposed and extensively developed by [Friedman and Tukey \(1974a\)](#), recently used by authors like [Guo et al. \(2006\)](#) in for the classification of microarray gene expression data. There is a vast literature on regularized discriminant analysis and regularized logistic regression, with a good number of the contributions dedicated to handling classification problems when $n \lll p$. It is important to note that it is quite typical to have contamination in the data whenever the dimensionality of the input space gets ever larger. The presence of outliers in the data is hard enough in low dimensional spaces ($n \lll p$), let alone in extremely high dimensional spaces where one now has to contend with both ill-posedness and outliers. Indeed, these situations trigger the need for both regularization (to deal with high dimensional ill posedness) and robustification to circumvent the ill-effect to outliers [Vanden Branden and Hubert \(2005\)](#). In the context of $n < p$, there is a relatively large literature on robust discriminant analysis with many of the contributions based on various approaches to robust estimation of both location and scatter [Filzmoser et al. \(2009\)](#), [Filzmoser and Todorov \(2011\)](#), [Filzmoser and Todorov \(2013\)](#), [Pires \(2003\)](#), [Pires \(2010\)](#), [Todorov and Pires \(2007\)](#), [Todorov and Filzmoser \(2009\)](#). Unfortunately, apart from [Vanden Branden and Hubert \(2005\)](#), [Pires \(2010\)](#), there has not been much work on robust discrimination when n is much less than p . In fact, we will reveal in our computational section that the traditional robust approach based on Minimum Covariance Determinant (MCD) estimation of the covariance structure fails miserably in the HDLSS context. The two approaches presented and explored by [Vanden Branden and Hubert \(2005\)](#) and [Pires \(2010\)](#) appear to still be in the very early stages of development. In our experimentations, we noticed that those techniques tend to work when n/p is close to 10^{-2} , but they all fail or struggle if the ratio n/p gets smaller.

In this paper, we explore both real life data - mainly microarray gene expression cancer data - and simulated data and we reveal patterns exhibited by the average test error as a function n , p and K (number of classes). In the context of simulated data, we also consider the impact of the contamination level/rate ϵ , the magnitude κ of contamination of the scatter matrix, and the level ρ of correlation among the predictor variables. Throughout our simulations, we use the same value for the contamination of the location.

The rest of this paper is organized as follows. In section two, we present a brief summary of discriminant analysis with an emphasis on where need for robustification and regularization arises. In section three, we discuss some of the most commonly used techniques of robustification, highlighting some of the limitations and merits of each method. In section four, we present the computational comparison of the techniques on real life data. In section five, we present a large simulation study, featuring various choices of the sample size n , dimensionality p , number of classes g , contamination level/rate ϵ and contamination size κ and correlation among the predictive variables ρ . We also highlight how various scenarios of these choices impact the average prediction error over R replications. In the section six, we present our conclusion and discussion along with a brief introduction to our future work dedicated to the regularized version of robust discriminant analysis in the HDLSS context.

2. Elements of Discriminant Analysis

Discriminant analysis is arguably one of the oldest and most commonly used approaches to pattern recognition. The main idea is that given K distinct groups one defines K different functions $\delta_k(\cdot)$, for $k = 1, \dots, K$, such that given a vector $\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)^\top$ whose i th row $\mathbf{x}_i^\top = (\mathbf{x}_{i1}, \mathbf{x}_{i2}, \dots, \mathbf{x}_{ip})$ for which class label \mathbf{y}_i is desired,

$$\hat{\mathbf{y}}_i = \widehat{class}(\mathbf{x}_i) = \arg \max_{k=1, \dots, K} \{\hat{\delta}_j(\mathbf{x}_i)\} \quad (2.1)$$

where $\hat{\delta}_k(\mathbf{x}_i)$ is the estimator of $\delta_k(\mathbf{x}_i)$. In other words, discriminant analysis estimates the class of a new object as the label whose discriminant functions yields the largest value given \mathbf{x}_i . Now, when the density function of \mathbf{x}_i given class k is known to be multivariate Gaussian, we have quadratic discrimination. If in addition all the classes have the same covariance matrix, i.e.,

$$(X|k) \sim \text{MVN}(\boldsymbol{\mu}_k, \Sigma) \quad (2.2)$$

we are in the presence of linear discriminant analysis (LDA). In other words, with LDA, each sample group has its class mean μ_k ($k = 1, 2, \dots, K$) but shares the sample covariance matrix Σ with the other groups. If $\pi_k = \Pr[Y = k]$ denotes the prior probability of class membership, then the LDA discriminant function is given by

$$\delta_k(\mathbf{x}) = -\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_k)^\top \Sigma^{-1}(\mathbf{x} - \boldsymbol{\mu}_k) + \log \pi_k. \quad (2.3)$$

In practice, $\delta_k(\cdot)$ is estimated from the data $(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_n, \mathbf{y}_n)$ by its empirical counterpart

$$\widehat{\delta}_k(\mathbf{x}) = -\frac{1}{2}(\mathbf{x} - \hat{\boldsymbol{\mu}}_k)^\top \hat{\Sigma}^{-1}(\mathbf{x} - \hat{\boldsymbol{\mu}}_k) + \log \hat{\pi}_k \quad (2.4)$$

where $\hat{\pi}_k = \frac{n_k}{n} = \frac{\sum_{i=1}^n w_{ik}}{n}$ is the observed proportion of group/class k observations and $\hat{\boldsymbol{\mu}}_k = \bar{\mathbf{x}}_k = \frac{1}{n_k} \sum_{i=1}^n w_{ik} \mathbf{x}_i$ is the empirical (sample) mean in class k , both defined using $n_k = \sum_{i=1}^n w_{ik}$, the number of observations from class k and the observed indicator of class membership given by

$$w_{ik} = 1_{\{\mathbf{y}_i = k\}} = \begin{cases} 1 & \text{if } \mathbf{y}_i = k \\ 0 & \text{otherwise} \end{cases} \quad (2.5)$$

Finally, the estimate $\hat{\Sigma}$ of the common covariance Σ is often the pooled covariance given by

$$\hat{\Sigma} = \frac{\sum_{k=1}^K n_k S_k}{\sum_{k=1}^K n_k - K} \quad (2.6)$$

where S_k is the sample (empirical) variance-covariance matrix of the k th class, namely

$$S_k = \frac{1}{n_k - 1} \sum_{i=1}^n w_{ik} (\mathbf{x}_i - \hat{\boldsymbol{\mu}}_k)^\top (\mathbf{x}_i - \hat{\boldsymbol{\mu}}_k). \quad (2.7)$$

It turns out that both the estimated location $\hat{\boldsymbol{\mu}}_k$ and the estimated scatter matrix $\hat{\Sigma}$ are sensitive to outliers, making the estimated discriminant function non robust under contamination. In the context of prediction, non-robustness leads to poor predictive performances due to the fact the presence of outliers in the explanatory variables causes the vital components of the discriminant function to be biased. We therefore need robust methods to estimate both location and scatter parameters in order to decrease the prediction error.

Additionally, when p is much larger than n , we encounter the extra problem of invertibility of $\hat{\Sigma}$. Typically one can solve this problem by selecting few subset of variables. However, more general approach deals with the problem by regularizing $\hat{\Sigma}$ using $\tilde{\Sigma} = \hat{\Sigma} + \lambda I_p$ or a convex version of it where $\tilde{\Sigma} = (1 - \alpha)\hat{\Sigma} + \alpha I_p$ with $\alpha \in [0, 1)$. Some of the earliest work on regularized discriminant analysis include the seminar paper by [Friedman and Tukey \(1974a\)](#), and later applications by [Guo et al. \(2006\)](#). Many authors have contributed extensively in the area of regularized discriminant analysis. Clearly, we both outliers and high dimensionality one needs to both robustify and than regularize. In the subsequent section we mainly focus on robustification having dealt with high dimensionality through projection pursuit [Pires \(2010\)](#), [Vanden Branden and Hubert \(2005\)](#).

3. Robust Estimation Methods for Linear Discriminant Analysis

In section two, one the challenges of discriminant analysis came from the fact that in the presence of outliers, the location and scatter estimates are not reliable because they are not robust. Many authors have contributed a wide variety of approaches to robustify LDA. One of the earliest approaches used to address this problem is a

technique known as Minimum Covariance Determinant (MCD) [Rousseeuw \(1984\)](#), [Rousseeuw \(1985\)](#). The literature on robust discriminant analysis has blossomed recently, with many papers studying and exploring various extensions of MCD. In this paper, we will be comparing the predictive performances of extensions developed ([Croux and Dehon \(2001\)](#)), then the one explored by [He and Fung \(2000\)](#) and [Hubert and Van Driessen \(2004\)](#). We will also look into the predictive performances of MCD extensions proposed ([Hubert and Van Driessen \(2004\)](#)) and the ones developed by [Todorov and Pires \(2007\)](#), [Hawkins and McLachlan \(1997\)](#). Another paper extending work on MCD is contributed by [Hubert and Debruyne \(2010\)](#). It is important to note all the above mentioned variations on the MCD theme have been implemented in R with packages like `rrcov` and `rrcovHD` readily available for immediate installation and use. Our goal in this paper is to consider different scenarios of contamination in high dimensional space and compare the performance of existing techniques, with the finality of establishing the conditions under each one of the techniques performs best.

[Rousseeuw \(1984\)](#)'s original Minimum Covariance Determinant (MCD) Estimator, can be briefly described as follows: given n observations $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n$ taken from a p -dimensional space $\mathcal{X} \subset \mathbb{R}^p$, with true location vector $\boldsymbol{\mu}$ and true scatter matrix $\boldsymbol{\Sigma}$, find the subset of h observations out of n such that the corresponding sample (estimated) covariance matrix yields the smallest determinant.

$$\det(\hat{\boldsymbol{\Sigma}}(\boldsymbol{\gamma}^{(\text{MCD})})) = \min_{\boldsymbol{\gamma} \in \{0,1\}^n} \left\{ \det(\hat{\boldsymbol{\Sigma}}(\boldsymbol{\gamma})) \right\}$$

where $\boldsymbol{\gamma} \in \{0,1\}^n$ simply represents that indicator vector such $\gamma_i = 1$ if observation i is among the h chosen $\gamma_i = 0$ if it is not. obviously, the final indicator vector $\boldsymbol{\gamma}^{(\text{MCD})}$ chosen by the MCD method is such that $\text{length}(\boldsymbol{\gamma}^{(\text{MCD})}) = |\boldsymbol{\gamma}^{(\text{MCD})}| = h$. Also, we use the notation $\hat{\theta}(\boldsymbol{\gamma})$ to denote the estimator of θ based on only the observations selected by the indicator vector $\boldsymbol{\gamma}$. Essentially, the MCD estimator of the location parameter $\boldsymbol{\mu}$ is defined by the mean of that subset $\boldsymbol{\gamma}^{(\text{MCD})}$ and the MCD estimator of the scatter matrix parameter $\boldsymbol{\Sigma}$ is defined by the covariance of that subset $\boldsymbol{\gamma}^{(\text{MCD})}$, more specifically

$$\hat{\boldsymbol{\mu}}_{\text{MCD}} = \hat{\boldsymbol{\mu}}(\boldsymbol{\gamma}^{(\text{MCD})}) \quad \text{and} \quad \hat{\boldsymbol{\Sigma}}_{\text{MCD}} = \hat{\boldsymbol{\Sigma}}(\boldsymbol{\gamma}^{(\text{MCD})}).$$

In practice, $\hat{\boldsymbol{\Sigma}}_{\text{MCD}}$ is chosen in such a way that it is a multiple of its covariance matrix. The multiplicative factor is chosen in such a way that $\hat{\boldsymbol{\Sigma}}_{\text{MCD}}$ is consistent at the multivariate normal model and unbiased for small samples ([Pison et al. \(2002\)](#)).

Now, it turns out that $\frac{n}{2} \leq h < n$ with $h = \lfloor (n+p+1)/2 \rfloor$ yielding the maximal breakdown point. Central to MCD is the fact that h must be determined or set. Also, it turns out that $\lfloor (n+h+1)/2 \rfloor < h < n$. In fact, it should be noted that the MCD estimator cannot be computed when $p > h$, since such a scenario would mean having a singular covariance matrix for any h -subset. It turns out that the whole MCD machinery needs $n \geq 2p$ in order to function at all. For our high dimension low sample size problems for which $p \gg n$, it is obvious that the basic formulation of MCD does not work. We discuss a little later the extension of MCD known as regularized MCD whereby the estimates of interest obtained in their regularized version.

Even within the satisfaction of the $n > 2p$ requirement, MCD is essentially very computationally intensive for the simple reason that the need to select a subset combinatorially to optimize a criterion requires a number of computing operations that can explode for even small sample sizes. Many fast versions of the basic MCD algorithm have been suggested, led by [Rousseeuw and Driessen \(1998\)](#).

It turns out that many variants of MCD lead to different predictive performances. There are several methods for estimating the common covariance matrix based on a high breakdown point estimator.

In the context of discriminant analysis, one of the obvious limitations of MCD lies in that fact that it trims all the classes/groups equally. Such an equal treatment of all the groups is potentially inefficient in the presence of uncontaminated groups. Many other drawbacks of the basic MCD approach to discriminant analysis have been scrutinized and addressed by authors such as [He and Fung \(2000\)](#), [Hubert and Van Driessen \(2004\)](#), [Christmann and Hable \(2013\)](#), [Croux and Dehon \(2001\)](#). Each of these variants was proposed in order to address/solve a perceived limitation/drawback of the basic MCD approach. Unfortunately, it turns out all these variants of MCD only work when p is less than n . In fact most of them require n to be at least greater than $2p$. In order to deal with n less than p situations one had to abandon MCD in the present form. As a matter of fact, we are currently working on an extension of the MCD approach that combines robustification and regularization to address HDLSS situation with n/p arbitrarily very small. In this paper however, we explore two of non MCD based techniques, namely robust SIMCA and projection pursuit (PP) discriminant analysis. As we will see in the computational section, PP will proved to be quite flexible but unfortunately fail when n/p gets to small

(less than to 10^{-2}).

SIMCA to High Dimensional Robust Classification: Soft Independent Modelling of Class Analogies (SIMCA) was introduced by Wold (1976). Biciato et al. (2003) explain that SIMCA ability to classify high dimensional data comes from the fact that it is based on a clever adaptation of principal component analysis (PCA). Thanks to the supervised nature of discriminant analysis (class membership known), SIMCA proceeds by performing principal component analysis in each of the K classes separately. Essentially, SIMCA can be summarized as an approach that combines robust PCA within each group based on robust covariance estimation to achieve good predictive performances in classification. More details on SIMCA can be found in Hubert and Van Driessen (2004), Hubert and Engelen (2004) Biciato et al. (2003) and Vanden Branden and Hubert (2005). SIMCA has been widely applied to areas as diverse as image analysis, microarray gene expression classification, and many other fields where data exists with n much less than p . An implementation of Robust SIMCA is provided through the R package `rrcovHD`, and will be used in our comparison of predictive performances of high dimensional robust classifiers.

Projection Pursuit Approach for Robust Linear Discriminant Analysis: From Wikipedia *Projection pursuit (PP) is a type of statistical technique which involves finding the most "interesting" possible projections in multidimensional data. Often, projections which deviate more from a normal distribution are considered to be more interesting. As each projection is found, the data are reduced by removing the component along that projection, and the process is repeated to find new projections; this is the "pursuit" aspect that motivated the technique known as matching pursuit. The idea of projection pursuit is to locate the projection or projections from high-dimensional space to low-dimensional space that reveal the most details about the structure of the data set. Once an interesting set of projections has been found, existing structures (clusters, surfaces, etc.) can be extracted and analyzed separately. Projection pursuit has been widely use for blind source separation, so it is very important in independent component analysis. Projection pursuit seek one projection at a time such that the extracted signal is as non-Gaussian as possible.* From the seminal article Friedman and Tukey (1974b), many authors like Friedman and Stuetzle (1981), Friedman (1987) have developed applications and extensions of Projection Pursuit to wide variety of statistical problems. There have also been many theoretical justifications and discussions of the strengths and appeal of Projection Pursuit Huber (1985), Hall (1990). Projection Pursuit has also been used in discriminant analysis and outlier detection Pan et al. (2000) and Polzehl and Forschungsgemeinschaft (1993). Pires (2010) presents one of the most recent developments on projection pursuit for robust discriminant analysis in high dimensional spaces. Her work builds first on the original projection pursuit idea as presented and developed by Friedman and Tukey (1974b) and Huber (1985), and combines robustness strategies and the foundational idea of projection pursuit presented in the above definition, to achieve robust linear discriminant analysis in high dimensional spaces. Four of the seven techniques compared in this paper are based on the projection pursuit approach to robustness linear discriminant analysis in high dimension spaces. As we'll later though, it will turn out that projection pursuit techniques will fail - somewhat catastrophically at times - when the intrinsic dimensionality of the data is high. This is unsurprising, since the very idea of PP presupposes the existence of lower dimensional space as the true/intrinsic basis of the data.

Ensemble Learning Approach to Robust Classification: It is often common in massive data that selecting a single model does not lead the optimal prediction. For instance, in the presence of multicollinearity which is almost inevitable when p is very large, function estimators are typically unstable and of large variance. This issue of inflated variance gets even more amplified in the presence of data contamination. It makes sense that when data contains outliers, they will cause the variance of estimators to increase. The now popular *bootstrap aggregating* also referred to as *bagging* offers one way to reduce the variance of the estimator by creating an aggregation of bootstrapped versions of the base estimator. This is an example of ensemble learning, with the aggregation/combination formed from equally weighted base learners. Consider regression for instance and let $\hat{g}^{(b)}(\cdot)$ be the b th bootstrap replication of the estimated base regression function $\hat{g}(\cdot)$. Then the *bagged* version of the estimator is given by

$$\hat{g}^{(\text{bagging})}(\mathbf{x}) = \frac{1}{B} \sum_{b=1}^B \hat{g}^{(b)}(\mathbf{x}).$$

If the base learner is a multiple linear regression model estimator $\hat{g}(\mathbf{x}) = \hat{\beta}_0 + \mathbf{x}^\top \hat{\beta}$, then the b th bootstrapped

replicate is $\hat{\mathbf{g}}^{(b)}(\mathbf{x}) = \hat{\beta}_0^{(b)} + \mathbf{x}^\top \hat{\beta}^{(b)}$, and the bagged version is

$$\hat{\mathbf{g}}^{(\text{bagging})}(\mathbf{x}) = \frac{1}{B} \sum_{b=1}^B \left(\hat{\beta}_0^{(b)} + \mathbf{x}^\top \hat{\beta}^{(b)} \right)$$

In this paper, our focus is on classification in high dimensional space where p is larger than n and the data also contains some outliers. Now, let's consider a multi-class classification task as defined much earlier with labels y coming from $\mathcal{Y} = \{1, 2, \dots, K\}$ and predictor variables $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_p)^\top$ coming from a p -dimensional space \mathcal{X} . Let $\hat{\mathbf{g}}^{(b)}(\cdot)$ be the b th bootstrap replication of the estimated base classifier $\hat{\mathbf{g}}(\cdot)$, such that $(\hat{y})^{(b)} = \hat{\mathbf{g}}^{(b)}(\mathbf{x})$ is the b th bootstrap estimated class of \mathbf{x} . The estimated response by bagging is obtained using the majority vote rule, which means the most frequent label throughout the B bootstrap replications. Namely, $\hat{\mathbf{g}}^{(\text{bagging})}(\mathbf{x}) = \text{Most frequent label in } \hat{\mathbf{C}}^{(B)}(\mathbf{x})$, where

$$\hat{\mathbf{C}}^{(B)}(\mathbf{x}) = \left\{ \hat{\mathbf{g}}^{(1)}(\mathbf{x}), \hat{\mathbf{g}}^{(2)}(\mathbf{x}), \dots, \hat{\mathbf{g}}^{(B)}(\mathbf{x}) \right\}.$$

Succinctly, we can write the bagged label of \mathbf{x} as

$$\hat{\mathbf{g}}^{(\text{bagging})}(\mathbf{x}) = \underset{y \in \mathcal{Y}}{\text{argmax}} \left\{ \text{freq}_{\hat{\mathbf{C}}^{(B)}(\mathbf{x})}(y) \right\} = \underset{y \in \mathcal{Y}}{\text{argmax}} \left\{ \sum_{b=1}^B (\mathbf{1}_{\{y = \hat{\mathbf{g}}^{(b)}(\mathbf{x})\}}) \right\}.$$

It must be emphasized that in general, ensembles do not assign equal weights to base learners in the aggregation.¹ For our purposes in this paper, it is crucial to address the issue of how bagging predictors can help achieve robust classification in high dimensional spaces. It turns out that that by combining bagging with subsetting of the input space, we end up with the so-called *Random Subspace Learning (Attribute Bagging)* which, as we'll argue later, yields very good predictive performances on high dimensional data contaminated with outliers. Consider the training set $\mathcal{D} = \{\mathbf{z}_i = (\mathbf{x}_i^\top, y_i)^\top, i = 1, \dots, n\}$, where $\mathbf{x}_i^\top = (\mathbf{x}_{i1}, \dots, \mathbf{x}_{ip})$ and $y_i \in \mathcal{Y}$ are realizations of two random variables X and Y respectively.

Suppose your goal is to use a total of M base learners to build an ensemble estimator $\hat{\mathbf{g}}^{\text{agg}}(\cdot)$ of the underlying function \mathbf{g} . Random Subspace Learning (Attribute Bagging) proceeds very much like bagging, with the added crucial step consisting of selecting a subset of the variables from the input space for training rather than building each base learners using all the p original variables.

- Randomly draw the number $d < p$ of variables to consider
- Draw without replacement the indices of d variables
- Build the d -dimensional model

This step is the main ingredient for variable importance estimation and also has the benefit of circumvent the limitation of bagging due to correlatedness of the trees making up the ensemble.

Let $\mathbf{z}_i \in \mathcal{D}$ be a random pair in the original training set \mathcal{D} of size n , and consider the bootstrapped sample $\mathcal{D}^{(b)}$ of size n also, generated by sampling with replacement from \mathcal{D} . Let $\Pr[\mathbf{z}_i \in \mathcal{D}^{(b)}]$ represent the proportion of observations from \mathcal{D} present in $\mathcal{D}^{(b)}$. Then it can be shown that $\Pr[\mathbf{z}_i \in \mathcal{D}^{(b)}] = 1 - \left(1 - \frac{1}{n}\right)^n$. As a result, if $\Pr[\mathbf{z}_i \notin \mathcal{D}^{(b)}] = \Pr[O_n]$ denotes the *proportion of observations from \mathcal{D} not present in $\mathcal{D}^{(b)}$* , then

$$\Pr[\mathbf{z}_i \notin \mathcal{D}^{(b)}] = \left(1 - \frac{1}{n}\right)^n = \Pr[O_n] \quad (3.1)$$

It turns out that $\Pr[O_n] \rightarrow e^{-1} \approx 0.37$ as $n \rightarrow \infty$. Which means that roughly about one third of the training set is not used when building the b th bootstrapped based learner. In the context of data contaminated with outliers, each observation - including outlier - has an asymptotic probability of e^{-1} of not affecting the base learner at hand. It is therefore reasonable to deduce that by averaging this exclusion of outliers over many replications (B base learners, with B typically large), one can achieve a robust classifier through bootstrap aggregation. Our conjecture in this regard is that through its bootstrapping mechanism, each outlier has a probability of e^{-1} of not affecting the base learner.

¹The general formulation in the context of regression for instance is $\hat{\mathbf{g}}^{(\text{agg})}(\mathbf{x}) = \sum_{b=1}^B \alpha^{(b)} \hat{\mathbf{g}}^{(b)}(\mathbf{x})$, where the aggregation is often convex, i.e. $\sum_{b=1}^B \alpha^{(b)} = 1$.

Algorithm 1 Random Forest via Random Subspace Learning

```

1: procedure RANDOMFOREST( $B$ )                                     ▷ The Random Forest Algorithm for  $B$  trees
2:   Choose a base learner  $\hat{g}(\cdot)$                                ▷ e.g.: Trees
3:   Choose an estimation method                                   ▷ e.g.: Recursive Partitioning
4:   for  $b = 1$  to  $B$  do
5:     Draw with replacement from  $\mathcal{D}$  a bootstrap sample  $\mathcal{D}^{(b)} = \{\mathbf{z}_1^{(b)}, \dots, \mathbf{z}_n^{(b)}\}$ 
6:     Draw without replacement from  $\{1, 2, \dots, p\}$  a subset  $\{i_1^{(b)}, \dots, i_d^{(b)}\}$  of  $d$  variables
7:     Drop unselected variables from  $\mathcal{D}^{(b)}$  so that  $\mathcal{D}_{\text{sub}}^{(b)}$  is  $d$  dimensional
8:     Build the  $b$ th base learner  $\hat{g}_{(b)}(\cdot)$  based on  $\mathcal{D}_{\text{sub}}^{(b)}$ 
9:   end for
10:  Use the ensemble  $\{\hat{g}_{(b)}(\cdot), b = 1, \dots, B\}$  to form the estimator

```

$$\hat{g}^{(\text{bagging})}(\mathbf{x}) = \frac{1}{B} \sum_{b=1}^B \hat{g}_{(b)}(\mathbf{x}).$$

```

11: end procedure

```

4. Comparison of Predictive Performances on real data

We now consider comparing the predictive performances of the techniques described earlier on real life data. Among other things, we present the apparent (training) error and the true (test) error which in this case is more precisely average test error over R replications as defined in (1.1).

$$\text{AVTE}(\hat{f}) = \frac{1}{R} \sum_{r=1}^R \left\{ \frac{1}{m} \sum_{j=1}^m l(\mathbf{y}_j^r, \hat{f}^{(r)}(\mathbf{x}_j^{(r)})) \right\}.$$

Throughout this paper, each replication randomly assigns 2/3 of the data to the training set and 1/3 to the test set. We do not consider a validation set because none of the techniques is based on a tuning parameter. We use $R = 200$ replications. We analyze 7 different datasets, six of which are high dimension low sample size (HDLSS) microarray gene expression datasets.

4.1. Description of the datasets

Diabetes: Our first data set deals with diabetes. It contains 145 observations and 3 variables. Three classes. This is obviously not a high dimensional dataset, but we used it here to reveal the stark difference in performance between methods when one switches from large n small p to large p small n . The data is available in the R package called diabetes mclust [Reaven and Miller \(1979\)](#).

Ceramic pottery data: This pottery data set was analysed by [Stern and Descoedres \(1977\)](#). Another authors have used it to test the robustness of their methods, namely [Cooper and Weekes \(1983\)](#) and later by [Pires \(2010\)](#).

These first two datasets are qualitatively different from all the other datasets explored in this paper. Indeed, all the other 5 datasets we explore here have in common the fact that they are all microarray gene expression data sets.

Prostate cancer data: Our second data set comes from microarray gene expression on prostate cancer data and is a subset of a data set from [Stephenson et al. \(2005\)](#)'s study. There were 37 samples which were classified as recurrent and 42 as non-recurrent primary prostate tumor.

Lymphoma dataset: The following data set is deals with Lymphoma and contains 180 observations and 661 variables.

Lung cancer dataset: The Lung cancer dataset analysed in this paper contains 197 observations and 1000 variables.

Colon cancer data: From [Alon et al. \(1999\)](#), it contains 62 observations on subjects classified into two groups (G1: subjects with colon cancer, with 40 observations; G2: healthy subjects, with 22 observations) and measured on 2000 variables (gene expression levels). The aim is to predict, as accurately as possible, the disease status from the gene expression levels. This is a well known data set in the modern classification literature (e.g., References from the paper) and the original version is available in the colonCA R package from Bioconductor. The raw

data is not normalized/preprocessed, which may lead to very bad classification results. Therefore a simple normalization procedure was applied: the data were log-transformed and after that each row was individually centered using its median.

Leukemia data: In this Leukemia data set, there are 3571 variables(features), 72 samples (Golub (1999)).

Brain cancer data: The last data set consider in this paper is a brain cancer dataset (Pomeroy (2002)). The total number of patients in this case is $n = 42$, each represented by $p = 5597$ microarray gene expression features, covering 5 different types of brain cancer.

Using the package `rrcov` and `rrcovHD` we computed the apparent error i.e. using the whole data

$\frac{n}{p}$	$\frac{145}{3} = 48.33$	$\frac{79}{500} = 0.158$	$\frac{180}{661} = 0.272$	$\frac{197}{1000} = 0.197$	$\frac{62}{2000} = 0.031$	$\frac{72}{3571} = 0.020$	$\frac{42}{5597} = 0.0075$
	Diabetes	Prostate	Lymphoma	Lung	Colon	Leukemia	Brain
Classic	13.10	NA	NA	NA	NA	NA	NA
Linda	10.35	NA	NA	NA	NA	NA	NA
PP	27.58	29.11	55.00	21.83	11.29	5.55	52.38
SIMCA	11.72	35.44	9.44	5.58	9.68	12.50	NA

TABLE 1

For the Variate of Dataset Mean probability of misclassification (Apparent Error rates) for the classical and robust estimators under different cases of sample size and covariates.

As we can see from the ; Linda (which means variance of the MCD approach) and classic LDA only work when n/p is greater than 1, in this case diabetes data. PP and SIMCA do handle HDLSS ($n/p < 1$) with the expectation SIMCA fail when $n/p < 10^{-2}$. When they both workfork both PP and SIMCA there is no clear winner. These seems to indicate that besides impact of n/p there is also the affect of the internal geometry of the data.

In the interest of comparing the predictive optimality of the techniques explored we now present the performances 4 different variants of PP discrimination, against robust SIMCA, RF and and diagonal discrimination. Here we use $R = 200$ replications and some of the results are summarized in table and graphs below.




FIG 1. Average test error as a function of n/p . The data sets appear on the x axis in decreasing order of n/p . The first one(diabetes) has $n/p = 145/3$ and the last one (leukemia) has $n/p = 72/3571$.

5. Comparison of Performances on Simulated Data


Throughout this paper, the multivariate Gaussian density will be denoted by $\phi_p(\mathbf{x}; \boldsymbol{\mu}, \Sigma)$

$$\phi_p(\mathbf{x}; \boldsymbol{\mu}, \Sigma) = \frac{1}{\sqrt{(2\pi)^p |\Sigma|}} \exp \left\{ -\frac{1}{2} (\mathbf{x} - \boldsymbol{\mu})^\top \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu}) \right\} \tag{5.1}$$

	PP-class	PP-huber	PP-mad	PP-sest	RSimca	RF	DDA
ceramic	8.44 (9.07)	9.44 (8.32)	12.50 (10.14)	11.06 (9.35)	23.00 (13.62)	16.00 (10.13)	9.72 (9.96)
colon-1	17.90 (7.73)	18.29 (7.99)	18.02 (7.68)	17.52 (8.07)	23.76 (10.34)	18.76 (8.96)	15.55 (6.78)
colon-2	23.38 (8.46)	23.05 (8.47)	21.67 (9.05)	21.19 (9.71)	25.88 (11.82)	21.71 (9.19)	31.86 (15.27)
diabetes	28.79 (5.41)	27.21 (5.68)	28.07 (5.57)	28.44 (5.40)	14.96 (4.95)	3.54 (2.53)	19.32 (4.99)
leukemia	5.58 (5.04)	5.21 (4.19)	5.13 (4.90)	4.85 (4.65)	15.56 (12.47)	3.04 (4.04)	3.21 (3.66)
lung	22.20 (4.18)	22.02 (4.18)	22.15 (4.28)	21.90 (4.09)	8.12 (3.37)	5.23 (2.53)	2.69 (1.69)
lymphoma	58.96 (6.02)	58.14 (5.95)	59.03 (6.36)	59.72 (5.63)	20.22 (7.29)	8.98 (3.98)	8.64 (3.55)

TABLE 2

Average test error along with the standard deviation.

FIG 2. Average test error on the lymphoma data set for which $n/p = 180/661$, and the number of classes is $g = 3$. These box plots compare the predictive performances of all the 7 methods considered.

Under the ϵ -contamination regime, the class conditional density of X in class k

$$p(\mathbf{x}|\boldsymbol{\mu}, \Sigma, k, \epsilon, \eta, \gamma) = (1 - \epsilon)\phi_p(\mathbf{x}; \mu_k, \Sigma) + \epsilon\phi_p(\mathbf{x}; \mu_k + \eta, \gamma\Sigma) \quad (5.2)$$

where η represents the contamination of the location while γ captures the level of contamination of the scatter matrix. Furthermore, in order to study the effect of the correlation pattern, we simulate the data using a covariance matrix Σ parameterized by τ and ρ and defined by

$$\Sigma = \Sigma(\tau, \rho) = \tau \begin{pmatrix} 1 & \rho & \cdots & \cdots & \rho \\ \rho & 1 & \rho & \cdots & \rho \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ \rho & \ddots & \rho & 1 & \rho \\ \rho & \cdots & \cdots & \rho & 1 \end{pmatrix} = \tau[(1 - \rho)\mathbf{I}_p + \rho\mathbf{1}_p\mathbf{1}_p^\top]$$

where \mathbf{I}_p is the p -dimensional identity matrix, while $\mathbf{1}_p$ is p -dimensional vector of ones. In fact, a more general covariance matrix to help better gauge the effect of correlation on the robust classification methods would be

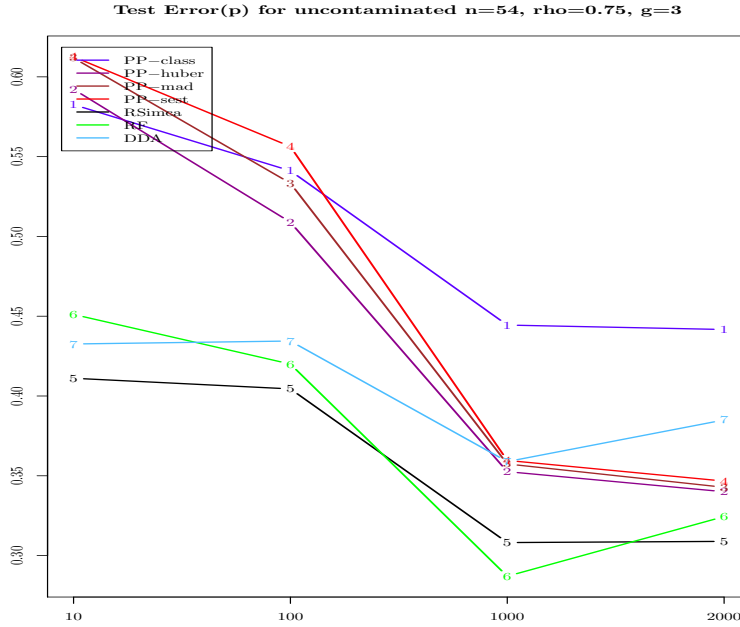


FIG 3. Average test error on the uncontaminated simulated data with $g = 3$ and $\rho = 0.75$. We herein reveal for each of the 7 methods, the effect of the input space dimension p on the average test error.

$\tau\Sigma$ where $\Sigma = (\sigma_{ij})$ with $\sigma_{ij} = \rho^{|i-j|}$.

$$\Sigma = \Sigma(\tau, \rho) = \tau \begin{pmatrix} 1 & \rho & \dots & \rho^{p-2} & \rho^{p-1} \\ \rho & 1 & \rho & \dots & \rho^{p-2} \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ \rho^{p-2} & \ddots & \rho & 1 & \rho \\ \rho^{p-1} & \rho^{p-2} & \dots & \rho & 1 \end{pmatrix}$$

For simplicity however, we use the first Σ with $\tau = 1$ throughout this paper. For the remaining parameters, we use $\epsilon \in \{0, 0.05, 0.15\}$, $\gamma \in \{9, 25, 100\}$, $K \in \{2, 3\}$ and $\rho \in \{0, 0.25, 0.75\}$ and $p \in \{10, 100, 1000\}$. As the vector of ϵ values shows, we consider 3 different levels of contamination, namely no contamination, mild contamination and strong contamination.

5.1. Uncontaminated Data

We first consider the performances of the techniques under an uncontaminated regime, i.e. $\epsilon = 0$. Our first simulation on under this regime looks at combination where the number of classes is $K = 2$ and then investigate the effect of ρ and p (input space dimension). As the plots all reveal, Projection Pursuit appears to perform very well (usually outperforming all the other methods) whenever intrinsic dimensionality of the data is low (captured by ρ very high) and the number of classes is 2 (Binary classification task).

	PP-class	PP-huber	PP-mad	PP-sest	RSimca	RF	DDA
10	58.26	59.22	61.19	61.26	41.11	45.11	43.26
100	54.11	50.89	53.33	55.63	40.44	42.00	43.44
1000	44.44	35.26	35.74	35.96	30.81	28.70	35.89
2000	44.17	34.00	34.28	34.67	30.89	32.44	38.50

TABLE 3

Average test error on the uncontaminated simulated data with $g = 3$ and $\rho = 0.75$. We herein reveal for each of the 7 methods, the effect of the input space dimension p on the average test error.

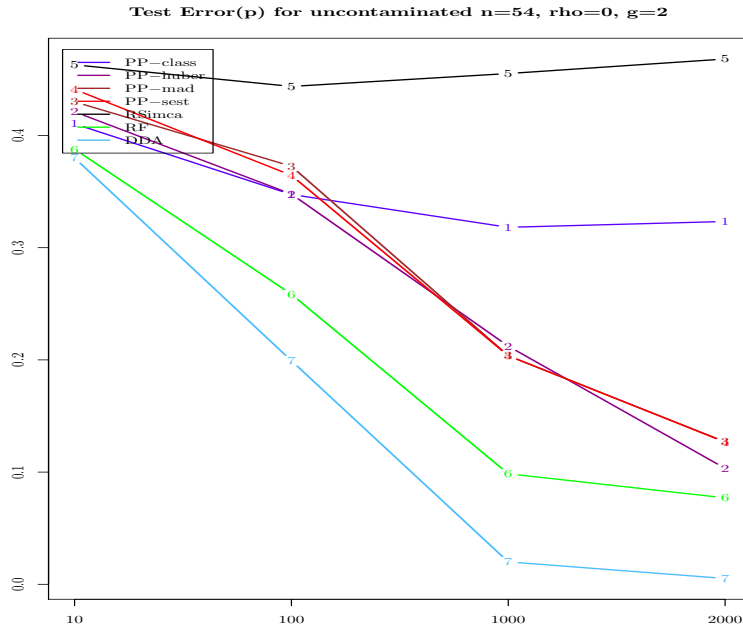


FIG 4. Average test error on the uncontaminated simulated data with $g = 2$ and $\rho = 0$. We herein reveal for each of the 7 methods, the effect of the input space dimension p on the average test error.

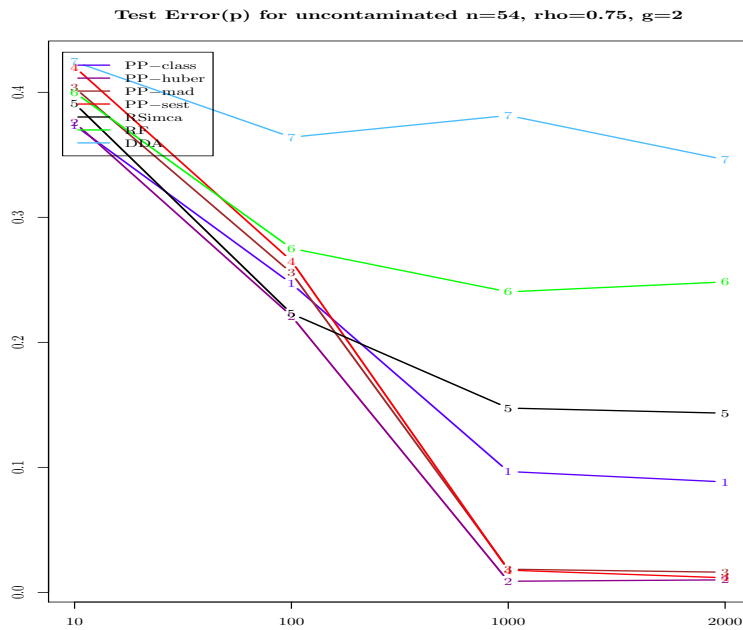


FIG 5. Average test error on the uncontaminated simulated data with $g = 2$ and $\rho = 0.75$. We herein reveal for each of the 7 methods, the effect of the input space dimension p on the average test error.

5.2. Effect of Mild Contamination

We now consider the performances of the techniques under a mildly contaminated regime, i.e. $\epsilon = 0.05$. Our first simulation on under this regime looks at combination where the number of classes is $g = 2$ and then investigate the effect of ρ and p (input space dimension) and κ .




FIG 6. Average test error on the mild contaminated simulated data with $g = 2$ and $\rho = 0.75$. We herein reveal for each of the 7 methods, the effect of the input space dimension p on the average test error.




FIG 7. Average test error on the mild contaminated simulated data with $g = 2$ and $\rho = 0.25$. We herein reveal for each of the 7 methods, the effect of the input space dimension p on the average test error.




FIG 8. Average test error on the mild contaminated simulated data with $g = 2$ and $\rho = 0.0$. We herein reveal for each of the 7 methods, the effect of the input space dimension p on the average test error.




FIG 9. Average test error on the mild contaminated simulated data with $g = 3$. We herein reveal for each of the 7 methods, the effect of the input space dimension p on the average test error.

5.3. Effect Strong Contamination

We now consider the performances of the techniques under a strongly contaminated regime, i.e. $\epsilon = 0.15$. Our first simulation on under this regime looks at combination where the number of classes is $g = 2$ and than investigate the effect of ρ and p (input space dimension) and κ .




FIG 10. Average test error on the strongly contaminated simulated data with $g = 2$ and $\rho = 0.75$. We herein reveal for each of the 7 methods, the effect of the input space dimension p on the average test error.




FIG 11. Average test error on the strongly contaminated simulated data with $g = 2$ and $\rho = 0.25$. We herein reveal for each of the 7 methods, the effect of the input space dimension p on the average test error.

6. Conclusion and Discussion

We have presented a thorough comparison of the predictive performances of several robust classification methods on high dimension low sample size data. On both real life and simulated data, interesting patterns emerged. We noted for instance that the SIMCA method, by being somewhat very general tends to yield mediocre predictive performances when p is much larger n , even though it rarely yield the worst among compared classification techniques. One of the most striking remarks in our study has to do with projection pursuit, the clearest being the fact that projection pursuit seems to do well only in binary classification. As a matter of fact, for all




FIG 12. Average test error on the strongly contaminated simulated data with $g = 2$ and $\rho = 0.00$. We herein reveal for each of the 7 methods, the effect of the input space dimension p on the average test error.




FIG 13. Average test error on the strong contaminated simulated data with $g = 3$. We herein reveal for each of the 7 methods, the effect of the input space dimension p on the average test error.

the scenarios involving more than two classes, projection pursuit seems to fail miserably regardless of all the other aspects of the data. Strikingly also, when there are only two classes, projection pursuit yields the best predictive performance if the correlation among the input space variables is large. This leads us to conclude that projection pursuit as a method for robust discriminant analysis is - at least in its present form - only best suited to binary classification for data whose intrinsic dimensionality is very low. For us, the most striking observation lies with the performance random forest. Indeed, as can be noted in all the computational results presented earlier, random forest tended to be the best. More precisely, there was no instance where random forest yielded the worst performance, and in most cases, it was either the very best or the second best. As we explained earlier, this can be explained by the very mechanism of random forest in the sense that at every iteration of the construction of a random forest, not only is the estimator based on a subset of input variables, but also crucially the bootstrap mechanism leaves out a proportion e^{-1} of the sample. This left out fraction certainly contains some of the outliers. It is our conjecture as indicated earlier, that the fact of leaving out a fraction of the data allows random forest to weed out outliers or at least average out their effect. Hence the inherent ability of random forest to achieve robustness by random subsampling. One could conjecture that the overall superior performance of random forest can be attributed to the fact it does both variable selection (by random subspace learning) thereby inherently addressing the extremely high dimensionality of the data, and also reduction (or even elimination) of the effect of outliers by subsampling. There is a sense here of a connection - however loose - between the subset selection of minimum covariance determinant (MCD) - recall that MCD select $h < n$ observations that yield the minimum covariance determinant - and the out-of-bag observations derived from the bootstrap in random forest. Therefore, we intend to investigate further the relationship between the fraction/proportion e^{-1} of random forest and the number h of observations used the MCD. We are also currently exploring various strategies of regularized MCD as a way to achieve robust classification in settings where n is much less than p .

Acknowledgements

Ernest Fokoué wishes to express his heartfelt gratitude and infinite thanks to Our Lady of Perpetual Help for Her ever-present support and guidance, especially for the uninterrupted flow of inspiration received through Her most powerful intercession.

References

- Alon, U., N. Barkai, D. Notterman, K. Gish, S. Ybarra, D. Mack, and A. Levine (1999). Broad patterns of gene expression revealed by clustering analysis of tumor and normal colon tissue probed by oligonucleotide arrays. *Proceedings of the National Academy of Sciences of the United States of America*, 96(12), 6745–6750.
- Bicciato, S., A. Luchini, and C. Di Bello (2003). Pca disjoint models for multiclass cancer analysis using gene expression data. *Bioinformatics*, 19(5), 571–578.
- Christmann, A. and R. Hable (2013, 25-30 August). On the bootstrap approach for support vector machines and related kernel based methods. In *Proceedings 59th ISI World Statistics Congress*, Hong Kong. Session CPS018.
- Cooper, R. and T. Weekes (1983). *Data, Models, and Statistical Analysis*. New York: Barnes and Noble,.
- Croux, C. and C. Dehon (2001). Robust linear discriminant analysis using s-estimators. *The Canadian Journal of Statistics*, 29(2), 473–493.
- Filzmoser, P., S. Serneels, R. Maronna, and P. Van Espen (2009). Robust multivariate methods in chemometrics. In B. Walczak, R. Ferre, and S. Brown (Eds.), *Comprehensive Chemometrics*, pp. 681–722.
- Filzmoser, P. and V. Todorov (2011). Review of robust multivariate statistical methods in high dimension. *Analytica Chimica Acta*, 705(1-2), 2–14.
- Filzmoser, P. and V. Todorov (2013). Robust tools for the imperfect world. *Information Sciences*, 245, 4–20.
- Friedman, J. H. (1987). Exploratory Projection Pursuit. *Journal of the American Statistical Association* 82(397), 249–266.
- Friedman, J. H. and W. Stuetzle (1981). Projection Pursuit Regression. *Journal of the American Statistical Association* 76(376), 817–823.
- Friedman, J. H. and J. Tukey (1974a). A projection pursuit algorithm for exploratory data analysis. In *IEEE TRANSACTIONS ON COMPUTERS*,.
- Friedman, J. H. and J. W. Tukey (1974b). A Projection Pursuit algorithm for exploratory data analysis. *IEEE Trans. on computers* 23(9), 881–890.
- Golub, T. e. a. (1999). Molecular classification of cancer: discovery and class prediction by gene expression monitoring. *Science*, 286, 531537. doi: 10.1126/science.286.5439.531.
- Guo, Y., T. Hastie, and R. Tibshirani (2006). Regularized linear discriminant analysis and its application in microarrays. *Biostatistics* 8(1), 86–100.
- Hall, P. (1990). On polynomial-based projection indices for exploratory projection pursuit. *The Annals of Statistics* 17(2), 589–605.
- Hawkins, D. M. and G. J. McLachlan (1997). High-breakdown linear discriminant analysis. *Journal of the American Statistical Association*, 92(437), 136–143.
- He, X. and W. Fung (2000). High breakdown estimation for multiple populations with applications to discriminant analysis. *Journal of Multivariate Analysis*, 72(2), 151–162.
- Huber, P. J. (1985). Projection Pursuit. *The Annals of Statistics* 13(2), 435–475.
- Hubert, M. and M. Debruyne (2010). Minimum covariance determinant. *WIREs Comp Stat*, 2: 3643. doi: 10.1002/wics.61 2, 3643.
- Hubert, M. and S. Engelen (2004). Robust pca and classification in biosciences. *Bioinformatics*, 20(11), 1728–1736.
- Hubert, M. and K. Van Driessen (2004). Fast and robust discriminant analysis. *Computational Statistics and Data Analysis*, 45(2), 301–320.
- Kondo, Y., M. Salibian-Barrera, and R. Zamar (2012). A robust and sparse k-means clustering algorithm. *arXiv preprint arXiv* (1201.6082).
- Pan, J.-X., W.-K. Fung, and K.-T. Fang (2000). Multiple outlier detection in multivariate data using projection pursuit techniques. *Journal of Statistical Planning and Inference* 83(1), 153–167.
- Pires, A. (2003). Robust linear discriminant analysis and the projection pursuit approach. In R. Dutter, P. Filzmoser, U. Gather, and P. J. Rousseeuw (Eds.), *Developments in Robust Statistics*, pp. 317–329. Physica-Verlag HD. ISBN: 978-3-642-63241-9.
- Pires, A. M. (2010). Projection-pursuit approach to robust linear discriminant analysis. *Journal Multivariate Analysis*, (101), 2464–2485.
- Pison, G., S. Van Aelst, and G. Willems (2002). Small sample corrections for lts and mcd. *Metrika*, 55(1-2), 111–123.
- Polzehl, J. and D. Forschungsgemeinschaft (1993). Projection pursuit discriminant analysis. *Computational*

- Statistics and Data Analysis 20*, 141–157.
- Pomeroy, S. e. a. (2002). Prediction of central nervous system embryonal tumor outcome based on gene expression. *Nature*, *415*, 436–442.
- Reaven, G. and R. Miller (1979). An attempt to define nature of chemical diabest using a multidimensional analysis. *Diabetologica*, *16*, 17–24.
- Rousseeuw, P. (1985). Multivariate estimation with high breakdown point. In W. Grossmann, G. Pflug, I. Vincze, and W. Wertz (Eds.), *In Mathematical Statistics and Applications*, Dordrecht. Reidel Publishing Company.
- Rousseeuw, P. J. (1984). Least median of squares regression. *Journal of the American Statistical Association*, *79*(388), 871–880.
- Rousseeuw, P. J. and K. V. Driessen (1998). A fast algorithm for the minimum covariance determinant estimator. *Technometrics* *41*, 212–223.
- Stephenson, A.J.and Smith, A., M. Kattan, J. Satagopan, V. Reuter, P. Scardino, and W. Gerald (2005). Integration of gene expression profiling and clinical variables to predict prostate carcinoma recurrence after radical prostatectomy. *Cancer*, *104*(2), 290–298.
- Stern, W. and J.-P. Descoedres (1977). X-ray fluorescence analysis of archaic greek pottery. *Archaeometry*, *19*, 73–86.
- Todorov, V. and P. Filzmoser (2009). An object-oriented framework for robust multivariate analysis. *Journal of Statistical Software*, *32*(3), 1–47.
- Todorov, V. and A. M. Pires (2007). Comparative performance of several robust linear discriminant analysis methods. *REVSTAT - Statistical Journal*, *5*(1), 63–83.
- Vanden Branden, K. and M. Hubert (2005). Robust classification in high dimensions based on the simca method. *Chemometrics and Intelligent Laboratory Systems*, *79*, 10–21.
- Wold, S. (1976). Pattern recognition by means of disjoint principal components models. *Pattern Recognition*, *8*(3), 127–139.